

History-sensitive versus future-sensitive approaches to security in distributed systems

Alejandro Hernandez and Flemming Nielson

Department of Informatics, Technical University of Denmark
{aher,nielson}@imm.dtu.dk

We consider the use of aspect-oriented techniques as a flexible way to deal with security policies in distributed systems. Recent work suggests to use aspects for analysing the future behaviour of programs and to make access control decisions based on this; this gives the flavour of dealing with information flow rather than mere access control. We show in this paper that it is beneficial to augment this approach with history-based components as is the traditional approach in reference monitor-based approaches to mandatory access control. Our developments are performed in an aspect-oriented coordination language aiming to describe the Bell-LaPadula policy as elegantly as possible. Furthermore, the resulting language has the capability of combining both history- and future-sensitive policies, providing even more flexibility and power.

1 Introduction

Distributed Systems are designed to manage large amounts of information, so they must be secured [9] to provide confidentiality for the information managed by them. The emerging Aspect-Oriented [1] field has been targeted to some security approaches [8]. Recently, a framework named **AspectKB** [7] has been proposed, with which is possible to model process calculi-like distributed systems and to capture security properties in a realistic way, attaching security policies to each location and then combining the relevant security policies when an interaction between locations takes place.

The way of expressing security policies in the **AspectKB** framework refers to the traditional non-distributed information-flow [11] style of assuring security, which statically analyses the possible behaviours of the system in order to avoid any potential misuse in the future. In **AspectKB**, this is exploited by making access control decisions dynamically, yet not considering any state of the locations but possibly some potential future behaviour.

In this paper, we shall consider a multilevel access control policy [6], the Bell-LaPadula model [2], and show some complications when trying to capture such a policy in a distributed framework in general, and in particular in a framework whose security policies focus on *looking to the future*, since such multilevel policies are better suited for past analysis of how the system reached its current state.

We then propose an extension to the **AspectKB** framework, allowing to express also policies that *look to the past*. We do this by adding the notion of a *localised state* to the locations modelled in the extended framework and allowing the security policies to access those states to make their decisions. With this, multilevel policies as the Bell-LaPadula policy can be easily captured, and we show how.

Since the original **AspectKB** framework was already intended to combine different security policies, with the extension done in this paper both policies that look to the past and policies that look to the future can be expressed and even combined. This not only benefits when trying to capture specific policies (such as the Bell-LaPadula one), but it also allows us to model every policy in its original way, providing more flexibility to the resulting extended framework.

Moreover, we shall argue that for some situations, expressing a policy in its original way could more precisely capture what is intended, and this insight would mean that our extended framework is more powerful as well. We shall start discussing this latter issue in the remainder of this Section (Subsection 1.1). In Section 2 we present a review of the Bell-LaPadula policy in its original formulation, and then we assess the challenges of adapting it to a distributed setting. Section 3 gives a brief review of the Logic used for dealing with the combination of policies. In Section 4 we present our extended framework, and show how to precisely capture the Bell-LaPadula policy. We also discuss why the resulting framework is more flexible than the existing one. In Section 5 we conclude.

1.1 Limitations of looking to the future

The framework we shall be dealing with throughout this paper is the formal language **AspectKB**. In that framework, which follows a process algebraic approach, the processes are modelled as actions taking place in specific locations, and interacting with other locations modelled as well. Furthermore, security policies can also be modelled, following an aspect-oriented manner. The policies can express their intentions by analysing the continuation (namely the process after the current action) of the involved processes, so that it is possible to know in advance what a process might do in the future. This reflects an information-flow style of providing security.

However, this information-flow style is not as adequate as it was for sequential programs. Indeed, since the only process that can be statically analysed is the one that continues after the current action, all the possible outcomes that may occur due to other processes could not be predicted. This means that, when deciding whether to allow the interaction to happen or not, it is necessary to look to the future of just one process, and this can lead to two possible ways of obtaining imprecise decisions, either *over-approximation* or *under-approximation*.

For understanding what over-approximation is, let us assume we *pessimistically* expect that a particular action done by a process could, because of other processes we do not know about, lead to an insecure state. In this situation we may disallow the process to execute that action, but in some cases there might be no other process performing anything that would lead to an insecure state.

For understanding what under-approximation is, let us assume we *optimistically* expect that a particular action done by a process will not lead to an insecure state because the very same process will not perform another related action that leads to such a state. In this situation we may allow the process to execute that action, but in some cases there might be some other process that makes the system reach some insecure state, due to some interactions that could have been avoided if the action was disallowed.

Let us discuss a simple example, without going into syntactic and semantic details, but still thinking about distributed processes and policies.

Let us think about a security policy where we have different security levels, and every location is assigned to some level. We do not want any information to be leaked from any security level to lower ones. Then, we should allow a process, running in a given location, to read data from another location, as long as the following two conditions are met: first, the other location, where the data is right now, is in a security level not higher than the one where the process is running; second, the process will not try, in the future, to write information to locations with security levels lower than the level of the location where the data is right now, since this writing may be influenced by the reading previously done.

Let us assume now a particular situation where we have 4 locations (say *A*, *B*, *C* and *D*), and 3 security levels (say 1, 2 and 3). Let us assume the security levels are ordered as their values in natural numbers ($3 > 2 > 1$). Let us assume that location *A* is in security level 1, locations *B* and *C* are both in security level 2, and location *D* is in security level 3. Figure 1 contains three cases of such a situation, showing

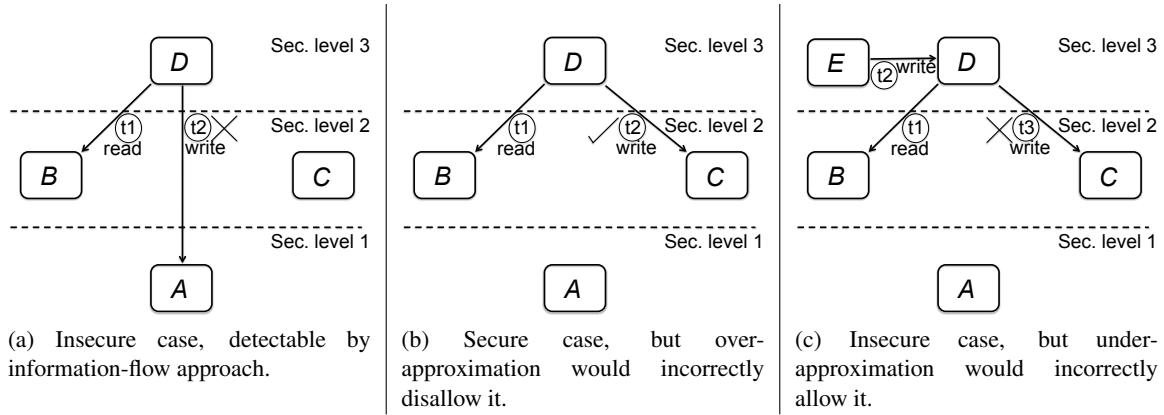


Figure 1: Examples of situations that might happen.

the locations and their security levels in different layers.

Illustrated in Figure 1a, there is a process in location *D*, that tries to read information from location *B* at *t1*, and then tries to write some information to location *A* at *t2*. This process should clearly be forbidden, because it does not meet the second condition of the policy we are trying to capture (although it meets the first one). Of course this can be done following the information-flow approach, looking to the future at *t1*, since we know that the process trying to read from *B* will try to write to *A*, and this should not be allowed.

However, let us think about another case, illustrated in Figure 1b. Let us say that the process running in location *D* whose first action is to read information from location *B* at *t1* then tries to write some information to location *C* (in the same level as *B*) at *t2*. This *does* meet the second condition of the policy, since the only information the process could write to *C* is what it has read from *B*. Therefore, this should be allowed.

But let us consider the next extension to the example, illustrated in Figure 1c. Assume there is a fifth location *E* that is in security level 3. Assume there is a process running in *E* that writes some information to *D* at *t2*, after the process running in *D* has read from *B* at *t1*. In this case, the future writing to *C* by the process running in *D* (which in this case will be done at *t3*) should be forbidden, because it might be influenced by the new information learned by location *D* at *t2*. Anyway, since the process that writes to *C* is not the same as the one running in *D*, the process algebraic way of modelling does not permit us to know in advance (at *t1*) that this will happen. If we had taken an approach of looking to the past, then we would have checked the insecure operation of writing to *C* right in the moment of the writing (at *t3*), and we would have known that some information from *E* would be leaked, and therefore we would avoid the write operation.

We could take the information-flow approach using over-approximation, and always avoid this type of write operation (e.g. from *D* to *C*, since the former is in level 3 while the latter in level 2), but that would be very imprecise (and restrictive), since sometimes there is nothing insecure in doing this write operation, as shown in the case of Figure 1b. Taking the information-flow approach using under-approximation would mean allowing the process in *D* to perform the read and the subsequent write, since that write operation is not insecure. This will be secure enough in the case of Figure 1b, but not in the case of Figure 1c.

So, we have found some possible situations where using an information-flow approach in a distributed setting is not completely precise, and therefore another approach might be taken, for instance

looking to the past. In the rest of this work, we shall be studying how to deal with looking to the past, and how to extend our distributed-systems framework to achieve this. We shall see that the resulting framework allows us to combine both approaches, therefore obtaining the advantages of both of them. In particular, we shall see that the simple example we have seen is just one possible instance of something that can easily (and more precisely) be captured by the Bell-LaPadula policy.

2 Assessment of the Bell-LaPadula model

In Subsection 1.1 we saw that for distributed systems the information-flow approach is not as adequate as it was for sequential programs. In this Section, we review another approach, the Bell-LaPadula (BLP for short) policy, and discuss the challenges of using it in a distributed setting, but aiming to show that this *can be* as adequate as in its original formulation.

2.1 The Operating System view of BLP

The BLP model is the most traditional Mandatory Access Control model. Here we briefly introduce it, inspired by [6], abstracting some unnecessary details that do not contribute to our study.

State. The computer system will be checked for security by looking into its state. For representing it, some sets must be introduced:

- S is the set of subjects that may use the information stored in the system,
- O is the set of objects (pieces of information) stored in the system,
- $A = \{\text{read}, \text{write}\}$ is the set of operations a subject may do over an object,
- L is a lattice of security levels.

Every state of the system is composed of a set of tuples of the form (s, o, a) (each tuple would mean that subject s is doing an a operation over object o), and of a tuple of functions (f_s, f_c, f_o) with types $S \rightarrow L$, $S \rightarrow L$ and $O \rightarrow L$. The functions are supposed to be total functions, and they will give, respectively, the maximum security level a subject can have (its *clearance*), the *current* security level a subject has¹, and the security level an object has (its *classification*). Formally, a state $(B, F) \in \mathcal{B} \times \mathcal{F}$, where $F = (f_s, f_c, f_o)$, and where:

- $\mathcal{B} = \mathbb{P}(S \times O \times A)$
- $\mathcal{F} = (S \rightarrow L) \times (S \rightarrow L) \times (O \rightarrow L)$

Policies. The BLP model specifies two properties that every state should meet in order to be considered secure.

- **ss-property**². A state (B, F) satisfies this property iff $\forall (s, o, a) \in B : a = \text{read} \implies f_s(s) \geq f_o(o)$. This means that each object being read by a subject should be in a level not higher than the level the subject is able to reach, which is usually called *no read-up*.

¹A subject can log into the system with a lower security level than its corresponding clearance. Once it did so, that security level cannot be changed until it logs in again.

²For “simple security”.

- **\star -property³**. This property consists of two parts. A state (B, F) satisfies the first part (let us name it **\star -property.1**) of this property iff $\forall (s, o, a) \in B : a = \text{write} \implies f_o(o) \geq f_c(s)$. This means that each object being written by a subject should be in a level not lower than the level the subject is currently in, which is usually called *no write-down*. On the other side, a state (B, F) satisfies the second part (let us name it **\star -property.2**) of this property iff $\forall (s, o, a) \in B : a = \text{write} \implies [\forall (s, o', a') \in B : a' = \text{read} \implies f_o(o) \geq f_o(o')]$. This means that if a specific subject (note the use of the same s in both quantifications) is operating with many objects, some being read and some being written, then no object being read could be in a higher level than any object being written. This prevents the subject to read some high-level object and then write a low-level one.

A state is said to be *secure* if it satisfies both properties.

2.2 The challenges of distribution

The BLP model was originally meant for Operating Systems. These have a particular feature: they are centralised, meaning that a central controller (i.e. the Operating System) takes care of everything that happens in the system. In particular it can control (and in some cases restrict) the processes that try to access resources. Moreover, one key concept needed for checking BLP policy compliance is the *state*, and since Operating Systems have a centralised state, they can do the calculations for knowing whether the BLP policy is met or not.

Lack of central controller. In a distributed setting we do not have any central controller; many locations run in parallel and share information, but no location can know what other locations are doing. Therefore, once a location is allowed access to some resource, there is no way other locations can forbid it from doing whatever it wants with the resource. In particular, there is no notion of state, processes interact and synchronise, but no central entity knows what has happened in the whole system so far.

It should be clear that a distributed framework is not trivially able to meet security properties that were originally developed for simpler systems, as for example centralised ones or sequential programs. In the case of Information-Flow approach, we have seen some simple examples where we can lose precision. In the case of BLP, in the next Subsection we propose an extension that will help us to adapt the policy to a distributed setting.

2.3 Extending BLP

The original formulation of the BLP policy relies on three functions, two of which can be applied to every subject and one to every object. They can be computed by the Operating System every time an action is to be executed, to check whether the resulting state will still be secure, and then decide whether to allow the action or not. Here we propose an extension to their domains to have common signatures, since in a setting without a central controller we might want to call any of them with any possible entity of the system without distinguishing between objects and subjects. We also propose a fourth function which captures information about the past interactions for each entity. Later in the paper we will see that this latter function can be used to have a form of localised state.

³Read “*star* property”. In some formulations of the BLP model, this property only consists of the first part because the ss-property uses f_c instead of f_s , and then the second part is just a consequence. However, that kind of formulation is again too restrictive, since a subject cannot perform read operations in levels up to its clearance, but just up to the level it has logged in.

The types of the three existing functions are then changed to $S \cup O \rightarrow L$ for all of them, and their definitions are extended in a straightforward way as follows:

$$\forall o \in O, s \in S : f_S(o) = f_O(o) \wedge f_C(o) = f_O(o) \wedge f_O(s) = f_S(s)$$

We call the new function f_H since it keeps track of (a part of) the *history* of the system. When we apply this function to a particular input subject (resp. object) we should learn what kinds of interactions the subject (resp. object) has been involved in during the past, therefore the output of the function would be a kind of current state of the argument subject (resp. object). To capture this notion of state, the function will not be fixed once and for all, as the original three functions were. Indeed, the output of this function will be:

- For a particular subject: the least upper bound of the security levels of all the objects read by the subject so far.
- For a particular object: the least upper bound of the security levels of all the subjects that have written to the object so far.

Formally, this can be expressed as follows (assuming $(B, (f_S, f_C, f_O, f_H))$ to be some “virtual” global state that depends on the interactions that have happened and $(B', (f_S, f_C, f_O, f'_H))$ the next one):

$$\begin{aligned} \forall (s, o, a) \in B : \quad & ((a = \text{read} \implies f_H(s) \geq f_O(o)) \wedge \\ & (a = \text{write} \implies f_H(o) \geq f_C(s))) \quad \wedge \\ & \forall (s', o', a') \in B' : (f'_H(o') \geq f_H(o) \wedge f'_H(s') \geq f_H(s)) \end{aligned}$$

This means that every time an interaction takes place, changing the state from $(B, (f_S, f_C, f_O, f_H))$ to some $(B', (f_S, f_C, f_O, f'_H))$, the output of f'_H for some input may be higher than or equal to that of f_H . It can actually be higher depending on the values of the entities read/written, for keeping the resulting least upper bound we expect to have. Indeed, a very simple result tells us that for every set \mathcal{L} :

$$\sqcup(\mathcal{L}) = \sqcup\{\sqcup(\mathcal{L} \setminus \{a\}), a\} \quad (\forall a \in \mathcal{L}) \quad (1)$$

And we should also observe that $\sqcup(\emptyset) = \perp$.


We shall use these four functions to capture this extended version of BLP in a distributed setting.

3 A brief review of Belnap Logic

For granting access according to some security policy, the traditional boolean values (**t** and **f**) are enough: **t** grants while **f** denies access. However, for a distributed setting, where policies might be contradictory (or not sufficiently informative), those two values might not be enough. We shall consider an extension to the Boolean Logic proposed by Belnap [3], which has been used for combining security policies [4].

In this extension to the boolean logic, two more values are considered: \perp and \top (read “bottom” and “top”). The traditional **t** would mean “the policy accepts the interaction” whereas the traditional **f** would mean “the policy does not accept the interaction”. Since different locations might aim at different security properties, their policies could be contradictory or they may lack information about some particular interaction. These situations can be represented by the two extra values that we have: \perp meaning “no decision” and \top meaning “contradiction” or “conflict”.

With this set of values, which we will call here **Four** (i.e. $\mathbf{Four} = \{\perp, \mathbf{t}, \mathbf{f}, \top\}$), it is possible to extend the usual boolean operations (\wedge and \vee) and to define new ones (\otimes and \oplus). For obtaining that, the set **Four** is equipped with two partial orderings, say \leq_k and \leq_t , as shown in Figure 2.

Figure 2: The Belnap bilattice **Four**: \leq_k and \leq_t .

The usual boolean \wedge is extended as computing the greatest lower bound in the \leq_t lattice, and the usual \vee as computing the least upper bound (thereby obtaining the same results as in boolean logic if the operands belong to $\{\top, \bot\}$). Analogously, the new operators over **Four** can be defined as computing the greatest lower bound (the \otimes operator) and the least upper bound (the \oplus operator), both in the \leq_k lattice.⁴

The negation operator \neg is extended by leaving the two new values unchanged (i.e. $\neg\bot = \bot$ and $\neg\top = \top$), and the implication \Rightarrow is extended as follows:

$$p_1 \Rightarrow p_2 = \begin{cases} p_2 & \text{if } p_1 \leq_k \top \\ \top & \text{otherwise} \end{cases} \quad \forall p_1, p_2 \in \mathbf{Four}$$

Another useful operator is the priority $>$, which returns the first operand unless it is \bot , in which case it returns the second operand. This would always consider what the first operand suggests unless it has no decision, in which case the second operand is considered.

4 Aspect-oriented framework for security

As mentioned, the **AspectKB** [7] framework allows us to express location-based systems in a process-calculus-oriented manner. This is achieved by extending the KLAIM [10] coordination [5] language. These *located processes* interact with other locations when they try to gather (or put) information from (or into) them (maybe themselves), which are usually named *tuple spaces*. The possibility of attaching to each location (regardless of whether it is a process or a tuple location) some security policy, which will govern the interactions the location may be involved in, turns the **AspectKB** language into an aspect-oriented language. Then, whenever an interaction takes place, the relevant policies are considered by the semantics to either grant or deny the interaction, using the four-valued Belnap Logic for deciding in a consistent way.

In this Section, an extension to that framework is made, mixing all process locations and tuple locations into just *entity locations*, and attaching to them more aspects than just the security policies. The extra information attached to each location refers to security levels in the sense of a multilevel security policy. Moreover, the mechanisms of the language explicitly keep track of some information (at a certain level of abstraction) regarding the interactions that have taken place, giving the flavour of a *localised*

⁴Notice that this could also be done by just extending the “truth tables” of the usual boolean operators and defining new ones for the new operators. That would mean, however, having not just 2 truth tables with 4 cells each (as in Boolean Logic) but 4 truth tables with 16 cells each, making it difficult to remember what each operator produces.

$N \in \mathbf{Net}$	$N ::= N_1 \parallel N_2 \mid l ::^w P \mid l ::^w \langle \vec{l} \rangle$
$P \in \mathbf{Proc}$	$P ::= P_1 \mid P_2 \mid \sum_i a_i.P_i \mid *P$
$a \in \mathbf{Act}$	$a ::= \mathbf{out}(\vec{\ell}) @ \ell \mid \mathbf{in}(\vec{\ell}^\lambda) @ \ell \mid \mathbf{read}(\vec{\ell}^\lambda) @ \ell$
$\ell, \ell^\lambda \in \mathbf{Loc}$	$\ell ::= u \mid l \qquad \ell^\lambda ::= \ell \mid !u$
$w \in \mathbf{Annot}$	$w ::= \langle lst, pol \rangle$
$lst \in \mathbf{LocSt}$	$lst ::= \langle \gamma^S, \gamma^C, \gamma^H, \gamma^O \rangle$
$\gamma \in L$	$left \text{ implicit}$

Table 1: **AspectKB+** Syntax – Nets, Processes, Actions and States.

state, which the semantics of the language keep updated⁵.

Besides this, one can write Aspects using the extra information, which is basically the output of the functions mentioned in Subsection 2.3 (considering that every entity location can be either a subject and/or an object in the whole system, so every location can be a potential input to all those functions). Then, this will allow us to capture, among others, the BLP policies without losing precision.

Following this informal introduction to our extension, which we shall call **AspectKB+** due to its enhanced features, we shall present its formalities.

4.1 Syntax and Semantics

Syntax. The syntax of the **AspectKB+** language is given in Tables 1 and 2. Table 1 gives the syntax for nets, the basic modules that can be described in the language. A net is a parallel composition of located processes and/or located tuples (data), together with an annotation explained below. Each process can be a parallel composition of processes, a non-deterministic choice between processes following an action, or a replicated process. A process not performing any action shall be written **0**. The allowed actions are reading from a location (**in** and **read**, resp. with or without deleting the data read) or writing to it (**out**).

Every location has an annotation w , whose first part (lst) is intended to keep track of the interactions that the location has been involved in. This *localised state* consists of 4 pieces of information (namely $\gamma^S, \gamma^C, \gamma^H$, and γ^O) that are elements of the lattice L of security levels (introduced in Subsection 2.1). Since every location can be input to the four functions f_S, f_C, f_H and f_O , and since the result of evaluating them is in L , we can keep attached to each location the result of evaluating each of those four functions.

The second part (pol) of the annotation in every location is the actual security policy governing the location, which has to be expressed using the syntax of Table 2. The policy can be a Belnap combination of policies, a boolean value, or a single aspect. This latter consists of a *cut* (the action, together with its continuation, to be trapped by the aspect), a condition *cond* (a boolean applicability condition) and a recommendation *rec* (a four-valued Belnap Logic advice for the aspect). To define an aspect, one may refer to the security levels stored in the trapped interaction or to a single value from the lattice L . To do the former, one can write an aspect naming some of the five syntactic names (S_s, C_s, H_s, O_t, H_t) specified in the category $v \in \mathbf{Lev}$, which will later be matched by the semantics to the specific values kept in the trapped interaction. To do the latter, one can write an aspect providing a specific value from L , as the category $v \in \mathbf{Lev}$ permits (by having γ among its choices). Finally, the **occurs-in** operator, which can be easily defined in a compositional way, checks whether the action occurs in the continuation process.

⁵As one can argue, having information inside the locations, namely the tuples, also gives us the flavour of state, yet that is information that changes according to what processes do, so we cannot rely on that information for guaranteeing any property.

$pol \in \mathbf{Pol}$	$pol ::= asp \mid \neg pol \mid pol \oplus pol \mid pol \otimes pol \mid pol \Rightarrow pol \mid$ $pol > pol \mid pol \wedge pol \mid pol \vee pol \mid \mathbf{true} \mid \mathbf{false}$
$asp \in \mathbf{Asp}$	$asp ::= [rec \text{ if } cut : cond]$
$cut \in \mathbf{Cut}$	$cut ::= \ell :: a^t . X$
$a^t \in \mathbf{Act}^t$	$a^t ::= \mathbf{out}(\vec{\ell}^t) @ \ell \mid \mathbf{in}(\vec{\ell}^{t\lambda}) @ \ell \mid \mathbf{read}(\vec{\ell}^{t\lambda}) @ \ell$
$rec \in \mathbf{Rec}$	$rec ::= \ell_1 = \ell_2 \mid \neg rec \mid rec \oplus rec \mid rec \otimes rec \mid rec \wedge rec \mid rec \vee rec \mid$ $rec \Rightarrow rec \mid \mathbf{true} \mid \mathbf{false} \mid a \text{ occurs-in } X \mid v_1 \geq v_2$
$cond \in \mathbf{Cond}$	$cond ::= \ell_1 = \ell_2 \mid \neg cond \mid cond_1 \wedge cond_2 \mid cond_1 \vee cond_2 \mid$ $\mathbf{true} \mid \mathbf{false} \mid a \text{ occurs-in } X$
$v \in \mathbf{Lev}$	$v ::= S_s \mid C_s \mid H_s \mid O_t \mid H_t \mid \gamma$ $\ell^t ::= \ell \mid _ \quad \ell^{t\lambda} ::= \ell \mid _$

Table 2: **AspectKB+** Syntax - Aspects for Security Policies.

Semantics. The semantics is given by a one-step reduction relation. It makes use of a structural congruence on nets (defined in Table 4), and also of an operation *match*, for matching input patterns to actual data, which could easily be defined in an inductive way by the structure of its arguments.

The reaction rules (defined in Table 3) prescribe how the system may evolve in the presence of some *process location* and some *target location*.

In the “where” lines of each rule, the boolean condition b is obtained by evaluating the security policies of the locations involved in the computation using the evaluation function $\llbracket \cdot \rrbracket$ (formally defined in Subsection 4.2). This is done to either allow or disallow the process to compute, and for this it also makes use of the function *grant* (also formally defined in Subsection 4.2) for turning four-valued Belnap truth values into boolean truth values. If the action was disallowed the involved process simply terminates, thereby becoming just a $\mathbf{0}$; otherwise the process evolves as the next paragraphs explain.

In the case of a **read** or **in** action, the process location l_s is subject to a substitution, using the result of the matching done with the *match* operation. Moreover, the localised state of that location might be modified, changing the historic component of its annotation by the least upper bound of the previous value and the security level of the target location l_t . This follows the suggestion of Equation (1).

In the case of an **out** action, the data is stored in the target location. However, this is not done directly, but actually another “virtual” location is created, with a special localised state. This is intended to permit the virtual location holding the pre-existing process Q to keep running as it was, without being interfered with. The virtual location now holding the data has an historic component on the annotation that is the least upper bound of the previous value in the location l_t and the security level of the process location l_s that has written the data. Of course it is possible that the value is the same as in the original l_t .

To simulate the log-in of a subject in a lower level than its clearance, a process can be annotated with a value for γ^C lower than the γ^S . This value will then never change, just as the γ^S and γ^O components of the localised state. Note also that the security policy annotating each location never changes either.

4.2 Meaning of policies and granting access

In the “where” lines of each semantic rule there is a check that tells whether the interaction should be allowed. For this purpose, the policies of both locations taking part in the interaction are combined using the Belnap operator \oplus , and the result of the evaluation by the operator $\llbracket \cdot \rrbracket$ is passed to the function *grant*.

The function *grant* is defined by $\text{grant}(p) = p \leq_k \mathbf{t}$, for all p in **Four**. The aim of granting access

$$\begin{array}{c}
\frac{N_1 \rightarrow N'_1}{N_1 \parallel N_2 \rightarrow N'_1 \parallel N_2} \qquad \frac{N \equiv M \quad M \rightarrow M' \quad M' \equiv N'}{N \rightarrow N'} \\
\\
(l_s ::^{w_s} \mathbf{read}(\vec{\ell}^\lambda) @ l_t.P + \dots) \parallel (l_t ::^{w_t} \langle \vec{l} \rangle) \\
\rightarrow \begin{cases} l_s ::^{w'_s} P\theta \parallel l_t ::^{w_t} \langle \vec{l} \rangle & \text{if } b \wedge \text{match}(\vec{\ell}^\lambda; \vec{l}) = \theta \\ l_s ::^{w_s} \mathbf{0} \parallel l_t ::^{w_t} \langle \vec{l} \rangle & \text{if } \neg b \end{cases} \\
\text{where } w_\delta = \langle \gamma_\delta^S, \gamma_\delta^C, \gamma_\delta^H, \gamma_\delta^O \rangle, \quad (\delta \in \{s, t\}); \\
\text{and where } b = \text{grant}(\llbracket pol_s \oplus pol_t \rrbracket (l_s :: \mathbf{read}(\vec{\ell}^\lambda) @ l_t.P, \langle \gamma_s^S, \gamma_s^C, \gamma_t^O, \gamma_s^H, \gamma_t^H \rangle)); \\
\text{and where } w'_s = \langle \gamma_s^S, \gamma_s^C, (\gamma_s^H \sqcup (\gamma_t^O \sqcup \gamma_t^H)), \gamma_s^O \rangle, pol_s \rangle. \\
\\
(l_s ::^{w_s} \mathbf{in}(\vec{\ell}^\lambda) @ l_t.P + \dots) \parallel (l_t ::^{w_t} \langle \vec{l} \rangle) \\
\rightarrow \begin{cases} l_s ::^{w'_s} P\theta & \text{if } b \wedge \text{match}(\vec{\ell}^\lambda; \vec{l}) = \theta \\ l_s ::^{w_s} \mathbf{0} \parallel l_t ::^{w_t} \langle \vec{l} \rangle & \text{if } \neg b \end{cases} \\
\text{where } w_\delta = \langle \gamma_\delta^S, \gamma_\delta^C, \gamma_\delta^H, \gamma_\delta^O \rangle, \quad (\delta \in \{s, t\}); \\
\text{and where } b = \text{grant}(\llbracket pol_s \oplus pol_t \rrbracket (l_s :: \mathbf{in}(\vec{\ell}^\lambda) @ l_t.P, \langle \gamma_s^S, \gamma_s^C, \gamma_t^O, \gamma_s^H, \gamma_t^H \rangle)); \\
\text{and where } w'_s = \langle \gamma_s^S, \gamma_s^C, (\gamma_s^H \sqcup (\gamma_t^O \sqcup \gamma_t^H)), \gamma_s^O \rangle, pol_s \rangle. \\
\\
(l_s ::^{w_s} \mathbf{out}(\vec{l}) @ l_t.P + \dots) \parallel (l_t ::^{w_t} Q) \\
\rightarrow \begin{cases} l_s ::^{w_s} P \parallel l_t ::^{w'_t} \langle \vec{l} \rangle \parallel l_t ::^{w_t} Q & \text{if } b \\ l_s ::^{w_s} \mathbf{0} \parallel l_t ::^{w_t} Q & \text{if } \neg b \end{cases} \\
\text{where } w_\delta = \langle \gamma_\delta^S, \gamma_\delta^C, \gamma_\delta^H, \gamma_\delta^O \rangle, \quad (\delta \in \{s, t\}); \\
\text{and where } b = \text{grant}(\llbracket pol_s \oplus pol_t \rrbracket (l_s :: \mathbf{out}(\vec{l}) @ l_t.P, \langle \gamma_s^S, \gamma_s^C, \gamma_t^O, \gamma_s^H, \gamma_t^H \rangle)); \\
\text{and where } w'_t = \langle \gamma_t^S, \gamma_t^C, (\gamma_t^H \sqcup (\gamma_s^O \sqcup \gamma_s^H)), \gamma_t^O \rangle, pol_t \rangle.
\end{array}$$

Table 3: Reaction Semantics of **AspectKB+**.

whenever the result is less than or equal to **tt** is for doing so not only if both policies agree with this, but also if some of the policies lack some decision, because this would mean that it does not actually forbid the interaction. This is related to the use of \oplus for combining the policies, and the aim is that whenever the policies are contradictory, the result of the evaluation by $\llbracket \cdot \rrbracket$ gives $\top \in \mathbf{Four}$, thereby denying access as long as at least one policy has evidence that the interaction should be disallowed⁶.

The evaluation function $\llbracket \cdot \rrbracket$ (Table 5) is defined inductively on the structure of the (infix) policy. The base cases are when the policy is just a constant (**true** or **false**) and when it is just an aspect (i.e. it belongs to **Asp**). In this latter case, the first (postfix) parameter, a *specific* action with continuation, is checked against the *cut* of the aspect, a *generic* action with continuation, using the function *check*, which could easily be defined in an inductive way by the structure of its arguments. This is achieved using a function *extract*, which produces the list of literals that occur in an action with continuation in a way that, for instance, $\text{extract}(\ell :: \mathbf{out}(\ell_1^t, \dots, \ell_n^t) @ \ell'.X) = [\ell, \mathbf{out}, \ell_1^t, \dots, \ell_n^t, \ell', X]$, which is done by just pattern matching the components of the given parameter and then pushing them into a list. The function *check* determines whether there is a substitution θ that can be performed in the *cut* that matches the parameter given to the $\llbracket \cdot \rrbracket$. This is needed because the *cut* can possibly consist of variables for representing the locations and even the arguments of the action in the *cut* may not be specified. If there is such θ , the *rec*

⁶This follows a conservative principle, as to actually grant access there should not be any policy at all denying the interaction.

$$\begin{array}{lcl}
l ::^w P_1 \mid P_2 & \equiv & l ::^w P_1 \parallel l ::^w P_2 \\
l ::^w *P & \equiv & l ::^w P \mid *P \\
l ::^w P & \equiv & l ::^w P \parallel l ::^w \mathbf{0}
\end{array}
\qquad
\frac{N_1 \equiv N_2}{N \parallel N_1 \equiv N \parallel N_2}$$

Table 4: Structural Congruence.

$$\begin{array}{l}
\llbracket [rec \text{ if } cut : cond] \rrbracket (l :: a.P, < \gamma_S, \gamma_C, \gamma_O, \gamma_{H_S}, \gamma_{H_t} >) = \\
\left(\begin{array}{l} \text{case } check(extract(cut); extract(l :: a.P)) \text{ of} \\ \text{fail : } \perp \\ \theta : \begin{cases} \llbracket ((rec \theta)\theta') \rrbracket & \text{if } \llbracket [cond \theta] \rrbracket \\ \perp & \text{if } \neg \llbracket [cond \theta] \rrbracket \end{cases} \\ \text{where } \theta' = [\gamma_S/S_S, \gamma_C/C_S, \gamma_O/O_t, \gamma_{H_S}/H_S, \gamma_{H_t}/H_t] \end{array} \right) \\
\llbracket [\neg pol] \rrbracket (N, \Gamma) = \neg(\llbracket [pol] \rrbracket (N, \Gamma)) \\
\llbracket [pol_1 \phi pol_2] \rrbracket (N, \Gamma) = (\llbracket [pol_1] \rrbracket (N, \Gamma)) \phi (\llbracket [pol_2] \rrbracket (N, \Gamma)), (\phi \in \{\oplus, \otimes, \Rightarrow, >, \wedge, \vee\}) \\
\llbracket [true] \rrbracket (N, \Gamma) = \mathbf{tt} \\
\llbracket [false] \rrbracket (N, \Gamma) = \mathbf{ff}
\end{array}$$

Table 5: Meaning of Policies in **Pol** for **AspectKB+**.

and the *cond* are substituted using it to determine the result. This is achieved using the usual two-valued meaning $\llbracket [cond] \rrbracket$, which could be straightforwardly adapted to a four-valued meaning $\llbracket [rec] \rrbracket$.

Due to the semantics of Table 3, the first parameter will always be the actual action taking place.

The second (postfix) parameter (consisting of five values in the lattice L) is used to produce another special substitution (θ') that is also used (together with θ) to determine the result of the recommendation *rec*. Due to the semantics of Table 3, the security levels annotated in the actual interacting locations are given here. Indeed, those taken from the target location are the ones identifying the *classification* (γ_t^O) of the location and the *historic* annotation (γ_t^H). Those taken from the process location are the ones identifying the *clearance* (γ_s^S) and the *current level* (γ_s^C) of the location and the *historic* annotation (γ_s^H).

It should be noticed that, while the θ substitutes according to some checking performed between the *cut* and the first parameter (the actual action), the θ' substitutes according to the five syntactic names prescribed by the syntax of Table 2, in the $v \in \mathbf{Lev}$ meta-variable. Therefore, when describing a system in **AspectKB+**, these syntactic names could be used to describe recommendations (*rec*) that will later be used to check actual security levels of the interacting locations, as already pointed in Subsection 4.1.

4.3 Capturing BLP in AspectKB+

Having developed our formal framework, we shall show how the extended BLP policy of Subsection 2.3 can be elegantly captured. We shall also show that we can easily decide which cases of the example in Subsection 1.1 are secure and which are not, without losing any precision, unlike the information-flow approach.

Remember that **AspectKB+** is a process calculus, and even though in the original formulation of BLP the compliance of a state with the policy is checked in *every* state, we can just check if a *transition* might take us to an “insecure state”. Also remember that **AspectKB+** provides us with the possibility, when describing aspects, of writing in the recommendation *rec* the five syntactic names we have mentioned,

which later will be substituted by the evaluation function $\llbracket \cdot \rrbracket$. So basically using those distinctive names we aim to capture the BLP policy.

The first aspects. Let us focus first on the **ss-property**, which prescribes that a subject cannot read an object that has higher security level than itself. The operations that can read information from other locations are the **read** and the **in** actions. So the aspects that capture the **ss-property** are the following:

$$\left[S_s \geq O_t \text{ if } l_s :: \text{read}(-)@l_t.P : \text{true} \right] \quad (2)$$

$$\left[S_s \geq O_t \text{ if } l_s :: \text{in}(-)@l_t.P : \text{true} \right] \quad (3)$$

Note that each aspect is trapping a particular operation, without caring about the parameters and with a trivial applicability condition. Whenever some of these aspects trap an action, the recommendation will be considered, granting access only if the security level of the subject is not lower than that of the object, since the two names S_s and O_t will then be replaced by the corresponding security levels of the actual interacting locations, thanks to Tables 3 and 5.

For the **★-property.1**, which prescribes that a subject cannot write any object that has lower security level than the level the subject is currently in, we have to follow a similar approach. Considering that the write operations are the **out** and the **in** (since deleting data is a form of write, because some implicit information could be communicated), the aspects are as follows:

$$\left[O_t \geq C_s \text{ if } l_s :: \text{out}(-)@l_t.P : \text{true} \right] \quad (4)$$

$$\left[O_t \geq C_s \text{ if } l_s :: \text{in}(-)@l_t.P : \text{true} \right] \quad (5)$$

Whenever some of these aspects trap an action, the recommendation will only grant access if the security level of the object is not lower than the one the subject is currently in (note the use of C_s instead of S_s).

The ★-property.2. Now let us consider the **★-property.2**, which was basically the one that initiated the proposal made in this paper, due to the difficulty of capturing it precisely in a distributed setting. Note that the semantics of **AspectKB+** will keep track of the least upper bound of the security levels of the objects read by a particular subject location, because it updates it whenever the subject reads something that is not lower than the current value. A similar observation can be done for the object locations.

Let us consider a subject location, which might have read some high information as long as its security level allows it (otherwise either aspect (2) or (3) would have denied it). Any subsequent write to a low location must be denied, and in principle either aspect (4) or (5) might decide this, unless the subject is logged into the system with a low security level. In any case, using the localised state that we have in the subject location, and making use of the H_s syntactic name provided by the syntax for expressing aspects, we define the following aspects:

$$\left[O_t \geq H_s \text{ if } l_s :: \text{out}(-)@l_t.P : \text{true} \right] \quad (6)$$

$$\left[O_t \geq H_s \text{ if } l_s :: \text{in}(-)@l_t.P : \text{true} \right] \quad (7)$$

They can be understood in a very similar way as aspects (4) and (5), with the difference being that they are considering the localised state of the subject location, instead of the level where the subject has logged into the system.

Analogous considerations can be done for an object location, and we can define the following aspects for finishing to capture the whole BLP policy:

$$\left[S_s \geq H_t \text{ if } l_s :: \mathbf{read}(-)@l_t.P : \mathbf{true} \right] \quad (8)$$

$$\left[S_s \geq H_t \text{ if } l_s :: \mathbf{in}(-)@l_t.P : \mathbf{true} \right] \quad (9)$$

Combining the aspects. After defining these eight aspects, the idea is to combine and attach them to every location, so every time an interaction is to take place, the semantics will consider all the aspects before allowing the interaction to happen.

Since the BLP model says that a state is secure if *both* properties are satisfied, then we need to make sure that none of the aspects representing the properties detects a possible insecure interaction, as that would mean that *at least one* of the properties is not satisfied. For capturing this situation, again the Belnap operator that must be used to combine the aspects for attaching them to the locations is \oplus .

Now we are ready to state our first Lemma:

Lemma 4.1 *If a distributed system is insecure in the sense of Section 2.3, then some of the aspects from (2) to (9) will deny the insecure interaction.*

For the converse we need to make an extra observation, discussed in the following Sub-subsection.

4.3.1 Initialising the historic value

The aspects just defined will check, among other values, the historic component γ^H attached to each location, and that value will be kept updated by the semantics. But, initially, one must give a particular value for the component. The chosen value will not affect the correctness of the aspects detecting insecure interactions, but to fulfil our requirement that we should not lose any precision while doing so (unlike the information-flow approach) the value should be $\perp \in L$. This follows the suggestion of Equation (1) and the observation just after it. Now we are ready to state our the converse of the previous Lemma:

Lemma 4.2 *If some of the aspects from (2) to (9) deny an interaction, then the hypothetical resulting global state, if the interaction was actually allowed, is insecure in the sense of Section 2.3.*

Furthermore, we can now easily verify that the three examples of Figure 1 are precisely captured. In particular, it should be taken into account what could happen after the process in location *E* writes to location *D* (Figure 1c). For the process in *D* to be actually influenced by this, it must explicitly read the data, since the semantics of **AspectKB+** will put it in another “virtual” location, with a higher historic component. So if the process is influenced, then at *t3* the aspects (actually aspect (6)) will prevent the write to *C*, otherwise the write will be allowed.

4.4 A very simple example

Let us now consider a very simple example to show how to combine looking to the future and to the past. Assume an airline has a database where information about the passengers is kept. The historic component of the database location is initialised to $\perp \in L$ so any process could read from it, but after

some data is written, only some processes could do so, according to the security level of the data written. The aspect that prescribes this is:

$$\left[\begin{array}{c} \text{clearance}_u \geq \text{history}_{\text{AirlineDB}} \\ \text{if } u :: \text{read}(\text{pass}, -)@ \text{AirlineDB}.P : \text{true} \end{array} \right] \quad (10)$$

As one can notice, this is a special case of aspect (8), but it is written like this here to emphasise the example.

One of the process locations that will not be allowed to read data from the database due to the previous aspect is the Government, whose clearance should not be enough to satisfy the *rec* of the aspect. Indeed, the historic component of the database should be high enough since the data written in there might be sensitive for the passengers.

However, in times of heightened security due to probable threats, the Government should be able to audit the passengers, therefore allowing it to read the database is necessary. Anyway, this should be allowed as long as the Government will not, later, give the passengers' data to the Press, to keep satisfying the right to privacy of the passengers. The following aspect prescribes this:

$$\left[\begin{array}{c} \neg(\text{out}(\text{data})@ \text{PressRelease} \text{ occurs-in } P) \\ \text{if } \text{Government} :: \text{read}(\text{pass}, \text{data})@ \text{AirlineDB}.P : \text{test}(\text{threatlevel}, \text{high})@ \text{AirlineDB} \end{array} \right] \quad (11)$$

This is just a *little* aspect that looks to the future⁷, where we see how this is achieved. In the presence of this aspect, the Government will be allowed to perform the read action, as long as there is a tuple $\langle \text{threatlevel}, \text{high} \rangle$ in the Airline database (i.e. the Airline was already notified of the heightened security situation), and also as long as the Government process trying to read the data will not leak the data to the Press in the future.

But one of the conditions is set in the *cond* of the aspect whereas the other in the *rec*. The reason is related to the fact that this aspect is a temporary one, and the aim is to combine it with the previous one. Moreover, the combination should be done in a way that the Government should actually be allowed to read the database, although the pre-existing aspect (aspect (10)) might deny this. Therefore, the operator needed for combining the two aspects is the priority $>$, and then the whole security policy for the Airline database would be $(11) > (10)$ ⁸.

With that, if the process location is the Government and the heightened security situation is declared, then aspect (11) will be considered. Otherwise, either the action will not be trapped by the aspect (if the process location is not the Government) or the condition *cond* will be **f** (if the threat level is not high), resulting in both cases in a $\perp \in \mathbf{Four}$ for aspect (11), considering then the aspect (10).

This example, even though it is very simple, clearly shows three features of our framework:

- The use of Aspects for security allows us to temporarily modify a distributed system without having to dig into the bussiness logic of the processes.
- The use of the four-valued Belnap Logic allows us to easily combine policies, providing flexibility for the aspect-oriented framework.
- The combination of looking to the past and to the future provides even more flexibility, giving the power to express exactly what is intended, for precisely satisfying some properties.

⁷In [12] there are many more realistic examples that look to the future, in the Electronic Health Records domain.

⁸Note that using this policy with the priority, the aspect (11) could even remain there, instead of just being a temporary one, since it will be ignored in most of the cases, as long as the tuple $\langle \text{threatlevel}, \text{high} \rangle$ is removed after the situation is normalised.

While the first two features were already present in the **AspectKB** framework (and in particular the first one is widely used in the aspect-orientation community), the third one is a very powerful add-on provided by the new **AspectKB+** framework.

5 Conclusion

We have studied the problem of enforcing multilevel security in a distributed system as precisely as possible. An information-flow approach poses the problem of having to “guess” what processes in other locations may do, thereby losing some precision. Therefore, we have extended an existing framework to deal with a notion of localised state, which has given us the power to access information about the past performance of the system, thereby allowing us to capture the Bell-LaPadula policy with precision.

The resulting framework provides a way to combine policies that look to both the future and the past due to the four-valued Belnap Logic. This gives flexibility to the framework, by capturing precisely what is intended by the security policies. This also gives more power than the previous framework of [7].

Acknowledgements. This work was partially funded by the Danish Strategic Research Council (project 2106-06-0028) “Aspects of Security for Citizens” and partially by the EU Integrated Project SENSORIA (contract 016004). We would like to thank Alan Mycroft for his comments on an early version of this paper. Finally, we really appreciated the comments from all the reviewers, they were very helpful.

References

- [1] G. Kiczales et al. (1997): *Aspect-oriented programming*. In: *ECOOP97, LNCS*, 1241, Springer, pp. 220–242.
- [2] D. E. Bell & L. J. LaPadula (1973): *Secure computer systems: mathematical foundations*. Technical Report, MITRE Corp.
- [3] N. D. Belnap (1977): *How a computer should think*. In: *Contemporary Aspects of Philosophy*, Oriel Press, pp. 30–56.
- [4] G. Bruns & M. Huth (2008): *Access-Control Policies via Belnap Logic: Effective and Efficient Composition and Analysis*. In: *CSF08*, IEEE Computer Society, pp. 163–176.
- [5] D. Gelernter & N. Carriero (1992): *Coordination Languages and their Significance*. *Communications of the ACM* 35(2), pp. 96–107.
- [6] D. Gollmann (1999): *Computer security*. Wiley.
- [7] C. Hankin, F. Nielson & H. Riis Nielson (2009): *Advice from Belnap Policies*. In: *CSF09*, IEEE Computer Society, pp. 234–247. Available at <http://doi.ieeecomputersociety.org/10.1109/CSF.2009.8>.
- [8] C. Hankin, F. Nielson, H. Riis Nielson & F. Yang (2008): *Advice for Coordination*. In: *COORDINATION08, LNCS*, 5052, Springer, pp. 153–168. Available at http://dx.doi.org/10.1007/978-3-540-68265-3_10.
- [9] J. McCune, S. Berger, R. Caceres, T. Jaeger & R. Sailer (2006): *Shamon: A System for Distributed Mandatory Access Control*. *ACSAC*.
- [10] R. De Nicola, G. Ferrari & R. Pugliese (1998): *Klaim: A Kernel Language for Agents Interaction and Mobility*. *IEEE Trans. on Soft. Engineering* 24(5), pp. 315–330.
- [11] A. Sabelfeld & A. C. Myers (2003): *Language-Based Information-Flow Security*. *IEEE J. Selected Areas in Communications* 21(1), pp. 5–19.
- [12] F. Yang, C. Hankin, F. Nielson & H. Riis Nielson (2010): *Aspects-oriented Access Control of Tuple Spaces*. Manuscript submitted to a journal.